

AMENDMENTS TO THE CLAIMS:

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

1. (currently amended) A method comprising the computer-implemented steps of:
gathering statistics by a database server about nodes that are stored in a database repository that is managed by the database server;
wherein said nodes form a hierarchy;
wherein each node is either an XML file or [[a]] an XML file container;
wherein at least one node in the hierarchy is an XML file container that contains a plurality of XML files, each of which contains a plurality of XML elements;
storing said statistics; and
in response to a request to the database server for access to one or more XML resources from said database repository, the database server computing a computational cost associated with each of two or more methods of accessing said one or more XML resources from said database repository, based on said statistics;
wherein the method is performed by one or more computing devices.
2. (currently amended) The method of Claim 1,
wherein the step of gathering statistics comprises gathering one or more data from a group consisting of:
a total number of nodes, in said hierarchy, that are accessible via a path through a specified node,
a total number of XML file containers, in said hierarchy, that are accessible via a path through said specified node,
a total number of nodes, in said hierarchy, that are accessible via a path through said specified node and that are in a level of said hierarchy that is immediately under a level of said specified node,

a total number of XML file containers, in said hierarchy, that are accessible via a path through said specified node and that are in a level of said hierarchy that is immediately under said level of said specified node, and a number of levels, from a root node of said hierarchy, at which said specified node is organized in said hierarchy.

3. (cancelled)
4. (original) The method of Claim 1, wherein the step of storing statistics comprises storing said statistics in a relational table of a database of which said database repository is part.
5. (previously presented) The method of Claim 4, wherein XML files of said nodes are XML resources, and wherein said relational table is a first relational table that is a different table than a second relational table in which said XML resources are stored in said database repository.
6. (original) The method of Claim 4, wherein said relational table is a relational table in which said XML resources are stored in said database repository.
7. (previously presented) The method of Claim 1, wherein XML files of said nodes are XML resources, and wherein the step of storing statistics comprises storing said statistics in a hierarchical index table in which said XML resources are indexed to said database repository.
8. (original) The method of Claim 1, wherein the step of computing a computational cost comprises computing a selectivity value for each of one or more predicates, from said request, that contain operators on said database repository.
9. (previously presented) The method of Claim 8, wherein XML files of said nodes are XML resources, and wherein each of said XML resources is stored, in association with a location of a node in said hierarchy, in a column of a table in said database repository, and wherein an operator contained in at least one of said one or more predicates is an operator that determines whether a particular XML resource can be located in said database repository through a particular specified path through a portion of said hierarchy.

10. (previously presented) The method of Claim 8, wherein XML files of said nodes are XML resources, and wherein each of said XML resources is stored, in association with a location of a node in said hierarchy, in a column of a table in said database repository, and wherein an operator contained in at least one of said one or more predicates is an operator that determines whether a particular XML resource can be located in said database repository at a terminal location of a particular specified path through a portion of said hierarchy.
11. (previously presented) The method of Claim 1, wherein XML files of said nodes are XML resources, and wherein the step of computing a computational cost comprises computing a computational cost of traversing, to locate a particular XML resource specified in said request, an index in which said XML resources are indexed to said database repository.
12. (original) The method of Claim 11, wherein computing said computational cost of traversing an index comprises computing a computational cost associated with one or more CPUs used for said traversing.
13. (original) The method of Claim 11, wherein computing said computational cost of traversing an index comprises computing a computational cost associated with reading data blocks in which portions of said index are stored.
14. (original) The method of Claim 11, wherein computing said computational cost of traversing an index comprises computing (a) a computational cost associated with one or more CPUs used for said traversing and (b) a computational cost associated with reading data blocks in which portions of said index are stored.
15. (previously presented) The method of Claim 1, wherein XML files of said nodes are XML resources, and wherein the step of computing a computational cost comprises (a) computing a selectivity value for each of one or more predicates, from said request, that contain operators on said database repository and (b) computing a computational cost of traversing, to locate a particular XML resource specified in said request, an index in which said XML resources are indexed to said database repository.

16. (original) The method of Claim 1, wherein said request for access to one or more XML resources from said database repository is a SQL query.
- 17.-20. (cancelled)
21. (original) The method of Claim 1, wherein said database repository is part of a relational database management system.
22. (previously presented) A computer-readable storage medium storing one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 1.
23. (previously presented) A computer-readable storage medium storing one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 2.
24. (cancelled)
25. (previously presented) A computer-readable storage medium storing one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 4.
26. (previously presented) A computer-readable storage medium storing one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 5.
27. (previously presented) A computer-readable storage medium storing one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 6.
28. (previously presented) A computer-readable storage medium storing one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 7.

29. (previously presented) A computer-readable storage medium storing one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 8.
30. (previously presented) A computer-readable storage medium storing one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 9.
31. (previously presented) A computer-readable storage medium storing one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 10.
32. (previously presented) A computer-readable storage medium storing one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 11.
33. (previously presented) A computer-readable storage medium storing one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 12.
34. (previously presented) A computer-readable storage medium storing one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 13.
35. (previously presented) A computer-readable storage medium storing one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 14.
36. (previously presented) A computer-readable storage medium storing one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 15.
37. (previously presented) A computer-readable storage medium storing one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 16.

38. (currently amended) A method comprising the computer-implemented steps of:
gathering, by a database management system, statistics about how many nodes that are
stored in a repository of said database management system satisfy certain criteria;
wherein said nodes form a hierarchy;
wherein each node is either an XML file or [[a]] an XML file container;
wherein at least one node in the hierarchy is an XML file container that contains a
plurality of XML files, each of which contains a plurality of XML elements;
wherein XML files of said nodes are XML resources;
storing said statistics in said database management system; and
the database management system using the statistics to determine how to process a query
that accesses one or more XML resources;
wherein the method is performed by one or more computing devices.
39. (previously presented) The method of Claim 38, wherein the step of storing comprises
storing said statistics as an XML data type in a schema-based table in said database
management system.
40. (previously presented) The method of Claim 38
wherein the step of gathering statistics comprises gathering each of
a total number of nodes, in said hierarchy, that are accessible via a path through a
specified node,
a total number of containers, in said hierarchy, that are accessible via a path
through said specified node,
a total number of nodes, in said hierarchy, that are accessible via a path through
said specified node and that are in a level of said hierarchy that is
immediately under a level of said specified node, and
a total number of containers, in said hierarchy, that are accessible via a path
through said specified node and that are in a level of said hierarchy that is
immediately under said level of said specified node.
41. (previously presented) A computer-readable storage medium storing one or more
sequences of instructions which, when executed by one or more processors, causes the
one or more processors to perform the method recited in Claim 38.

42. (currently amended) A method comprising the computer-implemented steps of:
in response to a request for access to one or more XML resources from a database repository within a database management system,
accessing, from said database management system, statistics about a structure of a hierarchy associated with said one or more XML resources;
wherein nodes form said hierarchy;
wherein each node of said hierarchy is either an XML file or [[a]] an XML file container; and
wherein at least one node in the hierarchy is an XML file container that contains a plurality of XML files, each of which contains a plurality of XML elements;
computing a computational cost associated with each of two or more methods of accessing said one or more XML resources from said database repository, based on said statistics;
wherein the method is performed by one or more computing devices.

43.-45. (cancelled)

46. (previously presented) A computer-readable storage medium storing one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 42.

47. (cancelled)

48. (currently amended) A system comprising:
one or more hardware processors;
means, executing on the one or more hardware processors, for gathering statistics by a database server about nodes that are stored in a database repository that is managed by the database server;
means, executing on the one or more hardware processors, for storing said statistics; and
means, executing on the one or more hardware processors, for computing, in response to a request to the database server for access to one or more XML resources from said database repository and based on said statistics, a computational cost, by the

database server, associated with each of two or more methods of accessing said one or more XML resources from said database repository;
wherein said nodes form a hierarchy;
wherein each node is either an XML file or [[a]] an XML file container; and
wherein at least one node in the hierarchy is an XML file container that contains a plurality of XML files, each of which contains a plurality of XML elements.

49. (currently amended) The method of Claim 1
wherein the step of gathering statistics comprises gathering statistics about at least one of:
(a) a median depth of a plurality of paths to a plurality of nodes in said hierarchy,
and (b) a maximum depth of a plurality of paths to a plurality of nodes in said hierarchy; and
wherein the plurality of nodes are accessible via a path through a specified node.
50. (cancelled).
51. (new) The method of Claim 1, wherein the request to the database server for access to one or more XML resources is through a view.
52. (new) The method of Claim 1, wherein the request includes one or more predicates and said one or more predicates includes at least one operator from the group of:
UNDER_PATH, and EQUALS_PATH.
53. (new) The method of Claim 8,
wherein a particular predicate of said one or more predicates includes an UNDER_PATH operator;
wherein the particular predicate is associated with (a) a depth of infinity, and (b) a particular node; and
wherein computing a selectivity value for the particular predicate further comprises:
determining a first number of nodes, in said hierarchy, that are accessible via a path through the particular node,
determining a second number of nodes, in said hierarchy, that are accessible via a path through a root node of said hierarchy,

dividing the first number of nodes by the second number of nodes to produce a third number, and
multiplying the third number by 100 to produce the selectivity value.

54. (new) The method of Claim 8,
wherein a particular predicate of said one or more predicates includes an UNDER_PATH operator;
wherein the particular predicate is associated with (a) a depth of one, and (b) a particular node; and
wherein computing a selectivity value for the particular predicate further comprises:
determining a first number of nodes, in said hierarchy, that are accessible via a path through the particular node and that are in a level of said hierarchy that is immediately under a level of said particular node,
determining a second number of nodes, in said hierarchy, that are accessible via a path through a root node of said hierarchy,
dividing the first number of nodes by the second number of nodes to produce a third number, and
multiplying the third number by 100 to produce the selectivity value.
55. (new) The method of Claim 8,
wherein a particular predicate of said one or more predicates includes an EQUALS_PATH operator; and
wherein computing a selectivity value for the particular predicate further comprises:
determining a particular number of nodes, in said hierarchy, that are accessible via a path through a root node of said hierarchy,
taking the inverse of the particular number of nodes, and
multiplying the inverse of the particular number of nodes by 100 to produce the selectivity value.
56. (new) The method of Claim 13,
wherein a particular predicate, from said request, includes an UNDER_PATH operator;
wherein the particular predicate is associated with (a) a depth of infinity, and (b) a particular node; and

wherein computing said computational cost associated with reading data blocks in which portions of said index are stored further comprises:

determining a number of XML file containers, in said hierarchy, that are accessible via a path through the particular node to produce said computational cost associated with reading data blocks in which portions of said index are stored.

57. (new) The method of Claim 13,

wherein a particular predicate, from said request, includes an UNDER_PATH operator; wherein the particular predicate is associated with (a) a depth of one, and (b) a particular node; and

wherein computing said computational cost associated with reading data blocks in which portions of said index are stored further comprises:

determining a number of XML file containers, in said hierarchy, that are accessible via a path through the particular node and that are in a level of said hierarchy that is immediately under said level of the particular node to produce said computational cost associated with reading data blocks in which portions of said index are stored.

58. (new) A database system comprising:

one or more hardware processors;

an XML data repository comprising XML files and XML file containers forming a hierarchy;

wherein at least one XML file container contains a plurality of XML files, each of which contains a plurality of XML elements; and

a database server, executing on the one or more hardware processors, that manages the XML data repository, wherein the database server is configured to:

gather statistics about the XML files and the XML file containers,

store said statistics,

receive a request for access to one or more XML resources from the XML database repository, and

compute a computational cost associated with each of two or more methods of accessing said one or more XML resources from the XML database repository, based on said statistics.

59. (new) A method comprising the computer-implemented steps of:
gathering statistics by a database server about XML files and XML file containers;
wherein the XML files and XML file containers are hierarchically stored in a database repository that is managed by the database server;
receiving a request to the database server for access, through a view, to one or more XML resources;
computing a selectivity value, based at least in part on the statistics, for a predicate included in the request; and
determining a query plan based, at least in part, on the selectivity value;
wherein the method is performed by one or more computing devices.
60. (new) A computer-readable storage medium storing one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 51.
61. (new) A computer-readable storage medium storing one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 59.